

AD-A131 320

NATURAL LANGUAGE COMMUNICATION WITH MACHINES: AN  
ONGOING GOAL(U) BOLT BERANEK AND NEWMAN INC CAMBRIDGE  
MA W A WOODS JUL 83 BBN-5375 N00014-77-C-0371

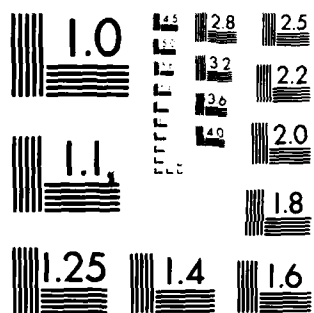
1/1

UNCLASSIFIED

F/G 6/4

NL


END  
DATE  
FILMED  
9 83  
DTIC



MICROCOPY RESOLUTION TEST CHART  
NATIONAL BUREAU OF STANDARDS-1963-A

**Bolt Beranek and Newman Inc.**

1770

12

Report No. 5375

**Natural Language Communication with Machines:  
An Ongoing Goal**

**Technical Report**

W.A. Woods

July 1983

Prepared for:  
The Office of Naval Research

**DTIC**  
**ELECTE**  
**S** **D**  
AUG 15 1983  
**D**

DTIC FILE COPY

**DISTRIBUTION STATEMENT A**

Approved for public release;  
Distribution Unlimited

**88 08 08 024**

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER BBN Report No. 5375	2. GOVT ACCESSION NO. AD-A131320	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle)  Natural Language Communication with Machines: An Ongoing Goal		5. TYPE OF REPORT & PERIOD COVERED  Technical Report
7. AUTHOR(s)  W. A. Woods		6. PERFORMING ORG. REPORT NUMBER BBN Report No. 5375
9. PERFORMING ORGANIZATION NAME AND ADDRESS Bolt Beranek and Newman Inc. 10 Moulton St. Cambridge, MA 02238		8. CONTRACT OR GRANT NUMBER(s)  N00014-77-C-0371
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research Department of the Navy Arlington, VA 22217		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE July 1983
		13. NUMBER OF PAGES 27
		15. SECURITY CLASS. (of this report)  Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Distribution of this document is unlimited. It may be released to the Clearinghouse, Department of Commerce, for sale to the general public.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)  Decision support, man-machine interface, natural language understanding, human-computer interaction, knowledge-based systems.		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This report is concerned with issues of man-machine interaction in decision support systems for high-level decision makers. It discusses components that such systems should have, what the current state of the art is with respect to such systems, and how current research in artificial intelligence is leading toward solving the remaining problems. Topics covered include natural language syntax and semantics, models of the beliefs and goals of the user, and knowledge-based helpful systems.		

DD FORM 1473  
1 JAN 73

EDITION OF 1 NOV 65 IS OBSOLETE

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Report No. 5375

**NATURAL LANGUAGE COMMUNICATION WITH MACHINES:**

**AN ONGOING GOAL**

**Technical Report**

July 1983

Principal Investigator:  
William A. Woods  
(617) 497-3361

**Prepared by:**

Bolt Beranek and Newman Inc.  
10 Moulton Street  
Cambridge, Massachusetts 02238

**Prepared for:**

The Office of Naval Research

Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A	



This research was supported by the Office of Naval Research under Contract No. N00014-77-C-0371. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Office of Naval Research or the U.S. Government.

TABLE OF CONTENTS

	Page
1. INTRODUCTION	1
2. KNOWLEDGE BASED LANGUAGE UNDERSTANDING FOR DECISION SUPPORT	3
3. THE NEED FOR FLUENCY AND CONCEPTUAL POWER	7
4. THE RATIONALE FOR NATURAL LANGUAGE UNDERSTANDING	9
5. BEYOND SYNTAX AND SEMANTICS	13
6. THE NEED FOR INTELLIGENT, HELPFUL SYSTEMS	15
7. THE NEED TO RECOGNIZE THE SPEAKER'S GOALS	17

8. THE NEED FOR KNOWLEDGE REPRESENTATION RESEARCH	21
9. STATUS AND PROSPECTS	23
10. CONCLUSION	25

## 1. INTRODUCTION

Suppose that you had five years in which to design a really good decision support system for high level decision makers. Where would you start? What would you try for? How would you do it? In this talk, I will try to give a sketch of the components that I think such a system should have, what the current state of the art is with respect to such systems, and how current research in artificial intelligence is leading toward solving the remaining problems.

I should begin by saying that what I have in mind by decision support is not a package of statistical decision procedures with respect to whose framework the decision maker is to express his options and valuations, after which the system will determine the optimal decision. Rather, I am concerned with the situations in which the decision maker has a problem, has not yet determined his options, much less his valuations, and is instead trying to come to understand the nature of the problem. Assuming that there is a computer system that contains extensive data among which the relevant information toward characterizing the problem may lie, how will the decision-maker find the relevant information and discover the patterns of information which will help him to understand what the problem is? How can a computerized system facilitate this task?

I am concerned with the unanticipated, nonstandard decision situations for which one cannot expect to have a predetermined package that displays just the right information in the right way to give the decision-maker what he needs. That is, for the situations (which I suspect are quite frequent) in which regular monthly or weekly reports do not give the whole picture, but



rather suggest questions that require further investigation, how do you construct a system that truly facilitates the investigative digging that is required to discover what is going on? How do you help a manager discover what is wrong when some aspect of the business is not going as expected?

A principal objective is to make the system sufficiently flexible that the decision maker can get information presented in whatever manner he finds helps him understand the situation, and to make it sufficiently intelligent and fluent that he can do this without having to take his attention away from the problem he is trying to solve and devote it instead to the issue of how to get the computer to do what he wants.

I will begin by describing a project at Bolt, Beranek and Newman that is addressing these issues. This project is attempting to make simultaneous, coordinated advances in a number of fundamental areas necessary for improving human communication with computers. These include fundamental techniques for the representation of conceptual knowledge, techniques for constructing helpful systems that can reason about a user's plans and goals, and techniques for efficiently parsing natural language and performing knowledge based inference.

## 2. KNOWLEDGE BASED LANGUAGE UNDERSTANDING FOR DECISION SUPPORT

BBN's research in Knowledge Representation and Natural Language Understanding is aimed at developing techniques for computer assistance to a decision maker in understanding a complex system or situation using natural language control of an intelligent graphics display. The motivating need is that of a military commander in a command and control context - especially in crisis situations. In such situations, not only does the commander need certain information in order to make his decisions effectively, but in complex situations, this requires the presentation of that information in a form that is matched to the abilities of human comprehension. A hypothetical scenario to illustrate the kinds of interaction we envisage is given below:

1. Cdr: Show me a display of the eastern Mediterranean.  
[computer produces display]
2. Cdr: Focus in more on Israel, and Jordan.  
[computer does so]
3. Cdr: Not that much; I want to be able to see Port Said and the island of Cyprus.  
[computer changes scale and window to include the desired features]
4. Cdr: Now show me the positions of all U.S. and Soviet vessels in the area.  
[computer does so, and makes a default assumption for displaying the difference between U.S. and Soviet vessels]
5. Cdr: Where is the John F. Kennedy?
- Computer: Two hundred miles to the west of the point displayed.

[The ship is not on the screen, so the system displays a point at the left edge of the display]

6. Cdr: Show me the course tracks for the Soviet vessels for the last five hours.

[computer does so]

7. Cdr: What kind of ship is that?

[points to a Soviet vessel]

Computer: Soviet missile cruiser.

8. Cdr: Show me the other missile cruisers, and display all vessel types with two digit code.

[computer blinks or flashes all of the missile cruisers for 2 1/2 seconds and displays with each vessel the two digit type code (assumed previously agreed on by the commander)]

9. Cdr: Remove the course tracks, and show small dots with one-hour course tracks for any known Soviet aircraft in the area.

[computer does so]

[commander makes his assessment of the situation and makes appropriate orders for his forces]

10. Cdr: Remove the planes and track the Soviet vessels for the next four hours. Show any deviations from current course double intensity and ring bell when detecting course change. Flash vessel changing course for 10 seconds.

[computer accepts standing orders for continual monitoring and conditional future behavior]

Notice that the user's utterances include imperatives to be taken as direct commands, as well as declaratives to be taken as indirect commands (as in exchange 3). Notice also, that the system is given considerable latitude to plan its response to

match what it "thinks" the user expects, rather than being meticulously instructed at a detailed level (for example, in choosing the exact boundaries of the region to display in response to utterance 1). In exchange 3, the user has simply stated the objectives that he wants to achieve at that point and left it to the system to determine how to achieve them.

We have conducted an experiment in collecting protocols of users interacting with simulated versions of the kinds of system we envision [4]. Our analysis of those protocols has convinced us that the behavior exhibited in exchange 3 above is the "tip of the iceberg" of a much more varied and common linguistic use. In particular, people often discuss a wide variety of changes they require in a system's response for reasons due both to changing their minds and to misunderstandings of what the system can or would do. Such users negotiate changes in particular ways, and they comment on the system's progress as the changes occur.

To explore the problems implicit in this scenario (which has not itself been implemented) and to develop techniques for dealing with them, we have been developing experimental prototypes in simplified domains. One of these is a system to support inspection and debugging of ATN grammars [1]. This system parses and interprets English requests, synchronized with pointing events on a screen, and produces appropriate display actions on a bit map graphics display in response. It permits a user to request portions of a display to be shown, objects in the display to be made visible or invisible, attributes of objects pointed to to be displayed, and specification of refocusing requests by means of statements of constraints on what is to be visible.

The system includes a sophisticated knowledge representation system, a comprehensive grammar of English, powerful general tools for natural language processing, and experimental capabilities for tracking the focus of attention in an ongoing communicative dialog, modeling the beliefs and goals of the user, recognizing the plan that underlies the user's utterances, and planning helpful responses to the perceived goals of the user.

A major accomplishment in this work has been the development of the knowledge representation system KL-ONE [1] and its use in the construction of the experimental prototype. KL-ONE is used to organize the semantic interpretation rules used to interpret sentences, to organize the models of the user's goals and beliefs (which are used to fill in details that are not explicit in the input), and to organize the knowledge of displays and display forms that are used to draw the pictures on the screen. The knowledge structuring capabilities of the KL-ONE system have proven themselves very powerful in this system, and the extent to which the same structures have proven useful in qualitatively different parts of the system gives evidence of the robustness of these capabilities.

Although the major thrust of our work is to address fundamental issues in a theoretically sound and general way (attempting to avoid the pitfalls of optimizing on aspects of particular applications), the results of the work to date have included not only increased understanding of the fundamental issues but also concrete subsystems that have been found useful by other groups in the scientific community.

### 3. THE NEED FOR FLUENCY AND CONCEPTUAL POWER

The underlying assumption of the BBN project is that in a crisis situation the commander needs an extremely flexible system, capable of manipulating large amounts of data and presenting it on a graphical display in a variety of ways until he feels satisfied that he has a grasp of the situation. Such a system would have abilities to display many kinds of map overlays, an ability to change the kinds and amounts of detail shown, an ability to conveniently construct unique displays to suit the situation at hand, as well as the ability to display tabular and graphical information and present textual material in ways that are easily comprehensible. This situation is not fundamentally different from the needs of a manager with a complex business decision to make.

In such circumstances, the display that the user wants and the modifications to it that he will subsequently want must be described in a highly fluent and expressive language, at a level of abstraction appropriate to the user's intent. One must not require the equivalent of a graphics programmer in order to obtain the displays required. Rather, one needs a system that is able to accept an abstract specification of the essential details of what should be in a display, and then intelligently and effectively determine the remaining details necessary to actually produce that display. This is true whether or not the actual specification of requests to the computer system is done by the decision-maker himself or by one or more subordinate specialists.

If the language of such a system is to be matched well to human cognitive abilities, it appears necessary for it to include a number of aspects of ordinary natural language, such as the use

of pronouns, the ability to take an incomplete specification and fill in the details on the basis of prior knowledge, and the ability to take a specification that would be potentially ambiguous out of context and determine the intended meaning. While artificial languages could perhaps be designed with the necessary properties, it is not obvious that one could do better than English as a language with sufficient power for expressing all of the needs of a manager in a complex decision task, while retaining the naturalness of use of English. Moreover, if one succeeded, it is likely that most of the computer processing difficulties inherent in understanding English would be present for this artificial language as well.

#### 4. THE RATIONALE FOR NATURAL LANGUAGE UNDERSTANDING

There are many advantages of natural language as a communication channel between a man and a machine. One of them is that the man already knows the natural language, so that he does not have to learn an artificial language nor bear the burden of remembering its conventions over periods of disuse. It also avoids his consciously translating (programming) his requests into the artificial language from the form in which they occur to him (presumably in a form very close to natural language). Especially for high-level personnel who use a computer system infrequently, or at least do not spend a major portion of their time dealing with the machine, these extra burdens of artificial language impose a severe barrier to the use of a machine.

Even for technical specialists who deal with a computer constantly, there is a distinction between the things that they do often and remember well, and many other things that may require consulting a manual and/or much conscious thought in order to determine the correct machine "incantation" to achieve the desired effect. For naive, inexperienced users, almost every transaction with the machine is of this form and the difficulty of deciding how to express a request is even more severe.

Whether a user is experienced or naive, and whether he is a frequent or occasional user, there arise occasions where he knows what he wants the machine to do and can express it in natural language, but does not know exactly how to express it to the machine. A facility for machine understanding of natural language can greatly facilitate the efficiency of expression in such situations -- both in speed and convenience, and in decreased likelihood of error.



A more important motivation for natural language understanding is the way that the underlying conceptual structures of English can match the user's conceptualization of the problem. Although most current natural language understanding systems do not achieve this goal, the understanding of the underlying English conceptual structure is far more important than the superficial resemblance to English syntax.

The problem of representation and use of conceptual knowledge in computers is one of critical importance in a wide variety of applications. These include not only intelligent, knowledge-based systems, but also general programming languages and systems. There is growing evidence that effective use of computers by both novices and experts and for both command interpreters and software development systems is greatly facilitated by structuring the computer program to use conceptual structures that correspond naturally to the conceptual structures that people use to organize the same information.

The issue of understanding the conceptual structure underlying natural language becomes especially important when the data to be manipulated by the machine is fundamentally natural language data. Such situations occur with interoffice memos, computer mail, and parts of requisitions, procurement specifications, etc. Current data management systems deal with data that can be fit into a relatively small number of predetermined formats, and they support requests of a relatively straightforward class. Forcing data into such formats usually leaves many things unexpressible in the data base, and the artificiality of the resulting data structure often makes the expression of many kinds of requests either impossible, or a difficult programming task.

It is my belief that an effective communication system for man-machine interaction in complex decision-making tasks must be essentially a natural language system. What is essential here is not the use of natural English words, although that has considerable mnemonic value, nor necessarily natural English constructions, although that considerably eases the learning task and the processing load required to use the system, but rather the natural English conceptual structure normally used for communication by humans. This conceptual structure has evolved through centuries of trial and error to become a very effective means of communication for an open ended set of complex ideas, situations, and goals. It is not only unlikely that we could design an artificial structure to match it, but such an artificial structure would also be difficult to learn and use.

Note that natural language does not preclude the introduction of abbreviations and telegraphic shorthands for complex or high frequency concepts -- the ability of natural English to accommodate such abbreviations is one of its strengths. Indeed, it is in the development of concise ways of saying relatively complicated things that natural language excels, and it is the development of a rich inventory of concepts in terms of which to express requests that would make such a system effective.

Bolt Beranek and Newman Inc.

Report No. 5375

## 5. BEYOND SYNTAX AND SEMANTICS

Natural language communication assumes a certain level of understanding (rather than mere decoding) on the listener's part. It is characterized by the use of such devices as pronominal reference, ellipsis, relative clause modification, natural quantification, adjectival and adverbial modification of concepts, and various attention focusing transformations. It is a vehicle for conveying concepts such as change, location, time, causality, purpose, etc. in natural ways. It also assumes that the understanding system has a certain awareness of discourse rules, enabling details to be omitted that can be easily inferred. It is the presence of such capabilities that marks what one should properly call a natural language understanding capability.

If the dialogues that I and others have considered can be taken as a fair sample, minimal requirements for a natural language system for use in decision-making applications include a facility for expressing quantification of actions or tests over sets of objects, for using adjectival and relative clause modification, for determining objects and sets of objects, for adverbial modification of verbs, for pronominal reference and definite noun phrase reference to objects or sets introduced previously in the discourse, for within-sentence pronominal reference, for conjunction and negation, for time and tense, and for extensive paraphrase variation in referring to objects and actions. In addition, it is highly desirable that the system be capable of understanding the various kinds of surface word-order transformations which people routinely apply (sometimes to eliminate ambiguity and sometimes to bring important aspects into

focus), and that it be capable of dealing with various kinds of ellipsis and vagueness.

Most of these areas have been studied by computational linguists and linguists, and many of them are sufficiently well understood that techniques for handling them are fairly well in hand. There is, however, a frequent tendency for claimed "solutions" to a given phenomenon to handle in truth only a restricted subset of the phenomenon. For some tasks, the restricted-case solutions are sufficient for useful application, but in general the details of a proposed solution need to be considered carefully to determine whether it is adequate for a particular application. Thus, one must be wary of assuming that a system that claims to handle pronominal reference and ellipsis (say) will actually handle either general examples of these phenomena or ones that actually arise in a particular application (even in the application for which the system was designed).

## 6. THE NEED FOR INTELLIGENT, HELPFUL SYSTEMS

Much of the time in communication with the system, the user will not say literally exactly what he means, and there are good reasons not to require him to do so. The major reason is that it is cognitively inefficient to be meticulously literal in one's communication (that's why computer programming is a time consuming and expensive activity). One of the major activities in programming a computer to do a complex task is the systematic specification of all of the details that would be left unsaid if one were instructing a human to carry out the same task. In the complex decision-aiding situations that we are considering, we cannot afford to require this degree of literal specification of detail. Rather, the system must know enough about the objectives of the user that it can fill in details in reasonable ways, asking the user for clarification occasionally, but only when absolutely necessary.

Moreover, the system should be able to use its general knowledge and the knowledge in its data base to go beyond merely doing what was requested, to provide additional information that can be inferred to be relevant to the user's goal and not otherwise known to the user. For example, when a military commander asks how many of his interdiction fighters are equipped with a particular kind of radar during a mission planning operation, the system should volunteer information about how many of those radars are out of commission (unless it knows that the commander already knows that). That is, the system should go beyond the passive execution of the user's commands to infer the goal structure underlying those commands where possible, and to volunteer additional relevant information (usually in accordance

to standing instructions as to what kinds of additional information should be offered in what situations).

Thus, in addition to understanding the syntax and semantics of language, a helpful decision system must share with the user:

1. Knowledge of the domain of discussion,
2. Knowledge of some of the user's intentions and goals,
3. Knowledge of what the user thinks the system can do.

This knowledge will make it possible for the user to interact with a flexible system that interprets his needs appropriately. For example, the user should be able to perform the following:

1. Request an action by the system or an effect to be accomplished where the level of description in the request is abstract and details are filled in by the system.
2. Ask questions whose proper interpretation depends explicitly or implicitly on the system's ability to infer some of the user's intentions.
3. Propose modifications of previous requests or of system responses where the system is to infer the relationship between the modification and the previous discourse.
4. Ask for clarifications, and then modify a request where the system provides help in response to the request for clarification and properly responds to the modified request.
5. Order the system to modify its overall future behavior where the system responds by changing its internal model of future action to conform with the order.

These features describe the kind of helpful system which I believe will be needed for complex decision-aiding tasks.

## 7. THE NEED TO RECOGNIZE THE SPEAKER'S GOALS

At this point one might ask whether this high level intentionality on the part of the system is really necessary, whether some more basic and simpler literal command interpreter might not be sufficient. There is in fact considerable evidence that when people attempt to solve problems with the aid of database management systems, they persist in expressing themselves to the machine as if it really understood what they were trying to achieve (in spite of overwhelming evidence to the contrary). Cohen, et al. [2] show that:

1. People, when asking questions of a database system, expect their intentions to be recognized and responded to.
2. People expect assistance from the question answerer in order to interpret the answers, correct their misconceptions of the underlying database and choose an alternate means to fulfilling their goals when a "dead end" is reached.
3. People do more than ask questions; they give commands, comments, clarifications and descriptions of their desires. Often these are expressed in terms of a previous system response.

When examined closely, the evidence indicates that such capabilities are essential for fluent communication. In fact, there are situations in which a user cannot even express his request without first coming to some understanding with the machine about the nature of the situation. Consider the following examples from actual user interactions with the PLANES system [5], a natural language question answering system that deals with Naval aircraft maintenance:



What was the average down-time of aircraft at each actorg? That includes maintenance time and AWM.

The user (who in this case was interfacing to the PLANES database through a person rather than a language understanding program) asked a simple question, but followed it with a comment which informed the system of what "counted" as down-time in interpreting the question. Interestingly, the PLANES database system could not take into account the user's purpose, so it printed a list of average down-times without indicating which actorg each time was associated with.

Another type of question users ask is evidenced by the example below. The question assumes that there were F14s on the Enterprise in March. The system, however, could answer "Zero" to this question even if no F14s were on the ship--the answer would be truthful but misleading, as it fails to take into account the user's intentions.

Of the F14 aircraft based on the Enterprise in March, how many were NOR while on the ship?

A third example illustrates how users and NL systems (actually persons simulating the part of a NL system) interact, making use of the discourse context, the intentions of each other, and the ability to modify a previous request.

U: How many cases of FOD were experienced

S: (THAT QUERY WOULD TAKE ABOUT 20 MINUTES  
TO COMPUTE) (IT CAN BE USEFULLY LIMITED BY  
CONSIDERING SPECIFIC AIRCRAFT)

U: Consider aircraft #27.

S: NIL

In this example, the system is aware of limits on the user's time and offers suggestions accordingly. The request to "Consider aircraft #27" is a vague request out of context, but here the user intends for the system to answer the question about FODs relative to the one aircraft.

A final example from the protocols of Genesereth [3] illustrates the implicit demands that people state in conversations. The user in this example is a MACSYMA user who is having difficulty solving equation D6 and has called on an advisor for help.

User: I was trying to solve D6 for Y, and I got 0.

Advisor: Did you expect COEFF to return the coefficient of D6?

User: Yes, doesn't it?

His remark to the advisor, which on the surface just explains his difficulty, also conveys his expectation that the advisor will explain why he's in trouble. The advisor's response question is understood to be part of his debugging assistance even though neither he nor the user state this explicitly.

The above examples illustrate the need for a system that can take into account the discourse context and the user's intentions in interpreting what the user wants the system to do. Constructing a system with these capabilities will require significant research in a number of areas, many of which have not been adequately studied. One of these is the need for situation dependent interpretation of linguistic devices such as "deixis" and "anaphora." The mechanism of anaphora permits one to make a subsequent reference to something that has previously been said in a dialog (e.g., using pronouns or definite noun phrases to

refer to previously mentioned objects), deixis involves similar references to things that have not been said but are present in some way in the nonlinguistic context of the conversation (e.g., in this case, what has just happened on the display screen). Anaphora has been extensively studied in linguistics (although the problems are far from solved), whereas deixis of the kind that occurs in the display context is considerably less well understood.

The resolution of both deictic and anaphoric reference requires a system to perform certain kinds of common sense inferences about the possible meanings of alternative possible referents, and to assess the plausibility of those alternatives. This in turn requires an ability to store and use considerable amounts of knowledge about the domain of discourse and the goals and objectives of the user. In addition to these linguistic devices, there is another level of interpretation of the user's input that depends even more critically on the use of such knowledge. This is the filling in of details that the commander can be assumed to have intended but did not literally say.

## 8. THE NEED FOR KNOWLEDGE REPRESENTATION RESEARCH

The above discussion illustrates the extent to which the representation and use of general world knowledge, knowledge of the domain, and knowledge of the goals and objectives of users are critical in the development of fluent communication and effective information display in the above context. Moreover, these problems are fundamental bottlenecks in a variety of other artificial intelligence applications. Consequently, a major portion of the BBN project has been devoted to fundamental problems of knowledge representation and use.

The KL-ONE knowledge representation system that we have developed during this project has an exceptionally good representation for the inheritance relations among structured concepts, including the relationships between corresponding parts of their structures. It has been used for representing a variety of different kinds of information in our current system, and has proven to be well structured in many respects. Some of the major features of KL-ONE are:

- o Inheritance of structured descriptions, i.e., when one concept is subordinate to another, the first "inherits" the properties of the second.
- o Taxonomic classification of generic knowledge, i.e., when a new generic concept is introduced into the network of existing concepts, it is automatically assimilated into the network at the right place so that it inherits the appropriate characteristics of related concepts and its properties are inherited by appropriate subordinates.
- o Intensional structures for functional roles, i.e., the use of a distinct class of nodes to formalize the notion of the role being played by an individual as distinct

from the individual itself. This distinction allows the specification of properties of roles when the filler of that role is unknown and allows the correct representation of the situation when a particular individual occupies several roles in an organization, for example.

- o Procedural attachment, i.e., associated with any concept or role one may record computational operations that are to be performed on instances of that concept. For example, a procedure for displaying an organization chart may be associated with the organization role of the concept of an organizational unit. Thus, when a user requests an organizational chart of a particular company, the procedure for producing the display can be inherited from the generic concept.

Knowledge representation research is one of the key problem areas in the development of intelligent natural language communication systems. There are many subtle problems of representation that are currently undergoing active investigation and whose solutions are required for really fluent natural language communication with machines. Some of the major ones have to do with the representation and reasoning about the goals and beliefs of other agents, modeling the structure of continuous discourse, modeling time and space, understanding what is required to represent and use abstract concepts and meta knowledge, and representing the subtleties of individual identity for entities that nevertheless change their properties over time.

## 9. STATUS AND PROSPECTS

The obvious question that one might ask about a natural language understanding capability of the kind I have described is "When can we have it?" If such capabilities do not now exist, when will they? These questions, unfortunately, do not have a crisp answer. On the one hand, limited natural language understanding systems do exist today. One of the earliest successful examples of such a system is the LUNAR system, which answers factual questions about the chemical analyses of the Apollo moon rocks [6]. On the other hand, there are aspects of natural language understanding that have not only not yet been solved, but whose solution has not yet been articulately envisioned. Thus, there is no point in the foreseeable future when the natural language problem will be totally solved.

For some applications, the existing level of capability demonstrated by the LUNAR system is adequate -- this includes the retrieval or computation of answers to specific factual questions of well-understood types from a data base of well-formatted information. It also includes the collection into a summary report of scattered pieces of data that satisfy a common predicate. However, there are many other capabilities that are required for a useful approximation to the kinds of capabilities that one would like. These include the resolution of pronominal reference (and other forms of anaphora), the resolution of ambiguity and/or vagueness in the user's requests, and the generation of helpful responses that take into account the perceived goals of the user. These capabilities are quite limited in current natural language understanding systems, and these are among the areas of primary research interest today.

Capabilities similar to LUNAR's are available today in systems such as Larry Harris's Intellect system and others that will be described in this conference. Moreover, there is likely to be a continued gradual incorporation of more discourse oriented capabilities into such systems. A major qualitative difference will occur, however, when systems become available that have a good conceptual model of the application domain with the same conceptual structure as that of the human user (as opposed to the data structure conceptualization typical of most data base management systems). The advent of such systems, however, will require further research in knowledge representation, knowledge based inference, common sense reasoning, and recognition of user's plans.

## 10. CONCLUSION

The need for natural language understanding systems arises in many different contexts, including computer-based decision making. Particularly in crisis situations, a decision maker must be able to interact with an information processing system as easily as with humans, and in a language that matches his own cognitive abilities. The system must be able to handle incomplete or inexact requests for information, resolve ambiguities, and use its knowledge of the domain of discourse and current context to understand pronomial and other indirect references. Although it might be possible to design an artificial language that would meet the above needs, the preferred method for developing such a capability is to use natural English as the communication language for such systems.

The primary advantages of English for such applications are:

1. The user need not learn and remember a large number of special conventions for communicating with the system;
2. The underlying conceptual structures of English can match the user's conceptualization of the task being performed;
3. One can use English to express instructions at varying levels of detail;
4. English includes shortcut devices such as anaphora (using "it", "the", etc. to refer to objects or phrases previously mentioned) and deixis (using non-linguistic devices, such as pointing to an object on a display screen) to specify one's intent.
5. The conceptual structure of English allows the incorporation of abbreviations and other shorthand devices for concisely expressing frequent commands.



Systems that allow natural language access to a conventional computerized data base are becoming available as software products on today's markets. Such systems generally have limitations in the range of English syntax that they will understand and severe limitations in their discourse understanding abilities. Moreover, they are dependent on the generally artificial conceptualizations of the domains that are built into their data bases. In the next few years, such systems can be expected to evolve somewhat more sophisticated discourse understanding abilities (better focus of attention models and handling of anaphoric reference), but will still fall short of intelligently understanding what the user wants and responding appropriately. Systems with the latter capabilities will emerge as further progress is made in knowledge representation, modeling belief systems, and common sense reasoning.

#### Acknowledgments

The research described here is the result of a large group effort involving several projects over a number of years. Participants in the group have included Madeleine Bates, Robert Bobrow, Ron Brachman, Phil Cohen, Brad Goodman, David Israel, James Schmolze, Candace Sidner, Marc Vilain, and Bonnie Webber. This research has been supported by ARPA under contract N00014-77-C-0378, monitored by ONR, and by ONR under contract N00014-77-C-0371.

# REFERENCES

- [1] Brachman, R.J., Bobrow, R.J., Cohen, P.R., Klovstad, J.W., Webber, B.L., Woods, W.A.  
Research in Natural Language Understanding - Annual Report: 1 Sept 78 - 31 Aug 79.  
BBN Report No. 4274, Bolt Beranek and Newman Inc., August, 1979.
- [2] Cohen, P.R., Perrault, C.R., Allen, J.  
Beyond Question Answering.  
In W. Lehnart and M. Ringle (editors), Strategies for Natural Language Processing, . Lawrence Erlbaum Assoc., Hillsdale, NJ, 1981.
- [3] Genesereth, M.R.  
Automated consultation for complex computer systems.  
PhD thesis, Department of Computer Science, Division of Applied Sciences, Harvard University, September, 1978.
- [4] Sidner, C.L.  
Protocols of Users Manipulating Visually Presented Information with Natural Language.  
Technical Report 5128, Bolt Beranek and Newman Inc., September, 1982.
- [5] Waltz, D.L., and Goodman, B.A.  
Writing a natural language data base system.  
In Proceedings of the Fifth International Conference on Artificial Intelligence, pages 144-150. International Joint Conference on Artificial Intelligence, Cambridge, MA, 1977.
- [6] Woods, W.A., Kaplan, R.M. and Nash-Webber, B.L.  
The Lunar Sciences Natural Language Information System: Final Report.  
BBN Report 2378, Bolt Beranek and Newman Inc., June, 1972.

Official Distribution List

Contract N00014-77-C-0371

	<u>Copies</u>
Defense Documentation Center Cameron Station Alexandria, VA 22314	12
Office of Naval Research Information Systems Program Code 437 Arlington, VA 22217	2
Office of Naval Research Code 200 Arlington, VA 22217	1
Office of Naval Research Code 455 Arlington, VA 22217	1
Office of Naval Research Code 458 Arlington, VA 22217	1
Office of Naval Research Branch Office, Boston 495 Summer Street Boston, MA 02210	1
Office of Naval Research Branch Office, Chicago 536 South Clark Street Chicago, IL 60605	1
Office of Naval Research Branch Office, Pasadena 1030 East Green Street Pasadena, CA 91106	1
Naval Research Laboratory Technical Information Division Code 2627 Washington, D.C. 20380	6

cont'd.

Naval Ocean Systems Center  
Advanced Software Technology Division  
Code 5200  
San Diego, CA 92152

1

Dr. A. L. Slafkosky  
Scientific Advisor  
Commandant of the Marine Corps  
(Code RD-1)  
Washington, D.C. 20380

1

Mr. E. H. Gleissner  
Naval Ship Research & Development Ctr.  
Computation & Mathematics Dept.  
Bethesda, MD 20084

1

Capt. Grace M. Hopper, USNR  
Naval Data Automation Command  
Code 00H  
Washington Navy Yard  
Washington, D.C. 20374

1

Mr. Paul M. Robinson, Jr.  
NAVDAC 33  
Washington Navy Yard  
Washington, D.C. 20374

1

Advanced Research Projects Agency  
Information Processing Techniques  
1400 Wilson Boulevard  
Arlington, VA 22209

1

Capt. Richard L. Martin, USN  
507 Breezy Point Crescent  
Norfolk, VA 23511

1

Director, National Security Agency  
Attn: R54, Mr. Page  
Fort G.G. Meade, MD 20755

1

Director, National Security Agency  
Attn: R54, Mr. Glick  
Fort G.G. Meade, MD 20755

1

Major James R. Kreer  
Chief, Information Sciences  
Dept. of the Air Force  
Air Force Office of Scientific Research  
European Office of Aerospace  
Research & Development  
Box 14  
FPO New York 09510

1

cont'd.

Mr. Fred M. Griffiee  
Technical Advisor C3 Division  
Marine Corps Development  
& Education Command  
Quantico, VA 22134

1

